arXiv:1601.03083v1 [cs.CG] 4 Jan 2016

ALGORITHMIC ASPECTS OF TVERBERG'S THEOREM

D. ROLNICK AND P. SOBERÓN

Abstract. We study the complexity of finding Tverberg partitions within various convexity spaces. For the classic geometric version of Tverberg's theorem, we obtain probabilistic algorithms that are weakly polynomial in both the dimension and the number of points. These algorithms extend to other variations, such as the integer version of Tverberg's theorem. For geodetic convexity on graphs, we show that the general problem of finding Radon partitions is NP-hard, and present efficient algorithms for certain special classes of graphs.

1. Introduction

Radon's lemma and Tverberg's theorem are central results in combinatorial geometry [24, 28]. These theorems describe the size at which a point set can be dissected into overlapping convex hulls.

Theorem (Tverberg's theorem)**.** *Given a set of* $(k-1)(d+1)+1$ *points in* $\mathbb{R}^d$*, there is a partition of the set into* k *parts such that the convex hulls of the parts intersect. Furthermore, this bound is optimal.*

The case $k = 2$ is known as Radon's lemma or Radon's theorem. There are many generalizations and extensions of Tverberg's theorem, such as colorful [1, 4, 5], topological [2, 3, 30], and quantitative versions [11] and are closely related to Helly's and Carathéodory's theorems. (See the survey [17] for a more in-depth presentation of the connections between theorems in combinatorial geometry.) The aim of this paper is to give new results regarding the algorithmic complexity of Tverberg-type results both in the classic geometric setting and in other convexity spaces.

We say a pair (X, conv) is a *convexity space* [29] if X is a set and $\text{conv} : 2^X \to 2^X$ is a function that satisfies

- $A \subset \text{conv}(A)$ for all $A \subset X$,
- if $A \subset B \subset X$, then $\text{conv}(A) \subset \text{conv}(B)$, and
- $\text{conv}(A) = \text{conv}(\text{conv}(A))$ for all $A \subset X$.

For instance, if $X = \mathbb{R}^d$ and $\text{conv}(\cdot)$ is the classic convex hull, we have a convexity space. However, the definition above is relaxed enough to capture many purely combinatorial instances. Given a connected graph G and two vertices $u, v \in V(G)$, we can define the segment $[u, v]$ as the union of all

Massachusetts Institute of Technology, Cambridge, MA. drolnick@mit.edu.
Northeastern University, Boston, MA. p.soberonbravo@neu.edu.

shortest paths from u to v. Then, for any $A \subset V(G)$, we can take $\text{conv}(A) = \cup_{u,v \in A}[u, v]$, yielding what is known as a *geodetic convexity space*. Tverberg-type results are interesting only for convexity spaces in which $\text{conv}(\emptyset) = \emptyset$.

Definition 1.1. *Given a convexity space* (X, conv)*, the* k*th* Radon number *(if it exists) is the smallest integer* r_k *such that for any* r_k *points in* X*, there is a* k*-partition such that the convex hulls of all parts intersect.*

Tverberg and Radon-type results in general convexity spaces are much more enigmatic than in the setting of $\mathbb{R}^d$. A classic conjecture by Eckhoff, known as the partition conjecture [18], states that if r_2 exists, then r_k exists and is at most $(r_2 - 1)(k - 1) + 1$. Thus, if this conjecture holds, Tverberg's theorem is a purely combinatorial consequence of Radon's lemma. Even though there have been convexity spaces proposed by Bukh showing that the partition conjecture does not hold in general [7], it is interesting to know for which convexity spaces it does hold.

It was shown by Duchet that the partition conjecture on finite convexity spaces in fact reduces to the case of geodetic convexity spaces on graphs [15]. Jamison corroborated the partition conjecture for certain graphs [19]. However, geodetic convexity spaces for general graphs are little understood, and even the Radon numbers remain unknown for many graphs. Results for grid graphs were recently presented by Dourado et al. in [13, 14].

For the classic notion of convexity in $\mathbb{R}^d$, it is not known if there exists an efficient algorithm for constructing Tverberg partitions under the assumptions of Tverberg's theorem. The case $k = 2$ (finding a Radon partition) is, however, simple; this problem reduces to identifying a linear dependence. A similar technique can also be applied to variations on Radon partitions, such as the colorful version of Radon's lemma [26], giving an algorithmic proof. The interesting challenge lies in Tverberg-type results with $k > 2$.

One may rewrite Tverberg's theorem as follows. Given a set S of n points in $\mathbb{R}^d$, there exists a partition into $m = \left\lceil \frac{n}{d+1} \right\rceil$ parts such that the convex hulls of the parts intersect. The problem of finding Tverberg partitions efficiently was motivated by the problem of finding centerpoints efficiently (see, for instance, [20]). Given a set S of n points in $\mathbb{R}^d$, we say a point p is an α-centerpoint of S if every closed halfspace containing p has at least αn points of S. Notice that any point in the intersection of a Tverberg partition $A_1, A_2, \ldots, A_m$ is immediately an (m/n)-centerpoint. As Tverberg partitions can be checked for correctness, this provides a robust way of finding centerpoints.

Obtaining optimal Tverberg partitions efficiently is out of reach for current algorithms. However, if one is willing to pay the price of decreasing the value of m slightly, better results can be obtained. For instance, there is a deterministic algorithm by Miller and Sheehy that solves the problem for $m = \left\lceil \frac{n}{(d+1)^2} \right\rceil$ in $n^{O(\log d)}$ time [20]. For faster running times in terms of n, there is an algorithm by Mulzer and Werner that solves the problem for

$m = \left\lceil \frac{n}{4(d+1)^3} \right\rceil$ in $d^{O(\log d)}n$ time [22]. Note that even though this last algorithm is linear in n, it is exponential in the dimension. Related algorithms can be found for variations of Tverberg such as the version with tolerance [21, 27]. In this variation, given a positive integer r, the goal is to find a Tverberg partition with the property that even after removing any r points the convex hulls of what is left on each part still intersect.

In Section 2, we show that one can go the other way around, using centerpoints to find Tverberg partitions efficiently. Our algorithms are not deterministic. They carry a probability ε of failure, which may be fixed in advance, as they depend heavily on the computation of approximated centerpoints by Clarkson et al. [9]. For $m \sim n/d^2$ we present an algorithm which is weakly polynomial in n, but exponential in the dimension. For $m \sim n/d^3$ we present an algorithm which is weakly polynomial in n, d and $\log(1/\varepsilon)$.

Theorem 1.2. *Let $m = \left\lceil \frac{n}{d(d+1)^2} \right\rceil$. There is an algorithm such that, given n points in $\mathbb{R}^d$, finds an m-Tverberg partition in time weakly polynomial in $n, d, \log(1/\varepsilon)$ with error probability ε.*

Theorem 1.3. *Let d, λ be fixed, where d is a positive integer and $0 < \lambda < \frac{1}{d+1}$, and $m = \left\lceil n\left(\frac{1}{d+1} - \lambda\right)\frac{1}{d}\right\rceil$. Then, there is an algorithm such that, given a set of n points in $\mathbb{R}^d$, finds an m-Tverberg partition in weakly polynomial time in n, $O(n^4 \log(1/\varepsilon))$, with error probability ε.*

In Theorem 1.3, if we allow d, λ to vary, we need an additional factor of $(d/\lambda)^{O(d)}$. Finally, we note an extension of our results to the integer setting. The following integral version of Tverberg's theorem is known to hold [16]:

Theorem. *Given k, d positive integers, there is an integer $T = T(k, d)$ such that for any set of T points in $\mathbb{R}^d$ with integer coordinates, there is a partition of the set into k parts such that the intersection of the convex hulls of the parts contains a point with integer coordinates.*

Note that even the case $k = 2$ is not completely solved, as the best current bounds are $5 \cdot 2^{d-2} + 1 \leq T(2, d) \leq d \cdot (2^d - 1) + 3$ and $T(k, d) \leq (k-1)d \cdot 2^d + 1$ [11, 23]. The algorithms we use can be extended to this setting, yielding the following:

Theorem 1.4. *Let d, λ, ε be fixed and that $0 < \lambda < 2^{-(d+2)}$. Then, there is an algorithm which is weakly polynomial in n such that, with error probability ε, for any set of n integer points in $\mathbb{R}^d$, it gives a partition of them into $\left\lceil \frac{n}{d}\left(\frac{1}{2^d} - \lambda\right)\right\rceil$ parts and an integer point $z \in \mathbb{R}^d$ which is in the convex hull of every part.*

In Section 3, we present algorithms for finding Tverberg partitions in geodetic convexity spaces induced by certain graphs. These algorithms rely on a similar notion of centerpoints to those we present for the classic convexity

setting, but are deterministic. Given a graph G, we say it is 2-separable if it is connected and for any two vertices u, v there are at most two edges such that removing them disconnects u from v.

Theorem 1.5. *Let G be a tree and $U \subset V$ be a subset of $2k$ vertices. Then, there exists a partition of U into k sets such that their convex hulls intersect; moreover this partition may be found efficiently.*

Theorem 1.6. *Let G be a 2-separable graph, and U a set of $4k-2$ vertices of G. Then, there exists a partition of U into k parts such that their convex hulls intersect; moreover this partition may be found efficiently.*

We conclude in Section 4 by presenting results on the hardness of finding Radon partitions for the geodetic convexity spaces induced by general graphs. We define the *graph-Radon problem* as follows:
Input: A graph $G = (V, E)$ and subset $W \subset V$.
Output: A decision whether or not there exists a partition (W^+, W^-) of W such that $\text{conv}(W^+) \cap \text{conv}(W^-)$ is nonempty.

Likewise, we define the *graph-Radon counting problem*:
Input: A graph $G = (V, E)$ and subset $W \subset V$.
Output: The number of partitions (W^+, W^-) of W such that $\text{conv}(W^+) \cap \text{conv}(W^-)$ is nonempty.

Theorem 1.7. *The graph-Radon problem is NP-hard and the graph-Radon counting problem is #P-hard.*

As a corollary to these results, we present a novel proof that it is NP-hard to compute the Radon number of the geodetic convexity space induced by a general graph, proven recently by Coelho et al. [10].

2. Algorithms for classic Tverberg partitions

We make use of Carathéodory's theorem in order to find Tverberg partitions.

Theorem (Carathéodory's theorem [8])**.** *Given a set $X \subset \mathbb{R}^d$ and a point $z \in \text{conv}(X)$, there is a subset C of at most $d+1$ points of X such that $z \in \text{conv}(X)$.*

Lemma 2.1. *Given X and z, we can find C using linear programming in time weakly polynomial in $|X|$ and d.*

Proof of Lemma 2.1. Deciding whether a point z is in the convex hull of a set Y can be expressed by the following linear program, where the λ_y represent variables and $y \in Y$ are d-dimensional constant vectors:

$$\sum_{y \in Y} \lambda_y y = z$$

$$\sum_{y \in Y} \lambda_y = 1$$

$$\lambda_y \geq 0 \text{ for all } y \in Y.$$

Linear programs can be solved in weakly polynomial time in the number n of variables and the number L of bits used to express the program. Specifically, Ye's interior point algorithm achieves a running time of $O(n^3L)$. We apply this observation as follows: For each x in X, we can check if z is in $\mathrm{conv}(X \setminus \{x\})$. If we succeed in finding such a point, we remove x from X and continue. Carathédory's theorem guarantees that we will be left with at most $d+1$ points containing z in their convex hull.

This procedure takes $\tilde{O}(n^4)$ time, with a tacit dependence on L; hence the algorithm is weakly polynomial. A slight improvement is achieved by the following procedure: Define a set of at most $d+1$ *essential* vertices in X, which contain z in their convex hull. If a subset of X does not contain z in its convex hull, then it must be missing at least one essential vertex. We can use this information to find the essential vertices by binary search: First eliminate half the vertices of X and check if their convex hull contains z. If it does contain z, continue to eliminate vertices; otherwise, replace half the vertices that were eliminated and try again. Finding each essential vertex in this fashion takes $\log n$ time, and there are at most $d+1$ of them. Therefore, we need $O(d \log n)$ applications of our LP-solver, giving total running time $\tilde{O}(n^3)$. We do not make an effort to optimize this running time further. □

Lemma 2.2. *Let α be in $(0,1]$ and d be a positive integer. Suppose there is an algorithm that runs in time $f(\alpha, d, n)$ such that given any set X of n points in $\mathbb{R}^d$ as input, it gives an α-centerpoint $z \in \mathbb{R}^d$ of X. In other words, every halfspace containing z contains at least αn points of X. Assume as well that given a set X of n points in $\mathbb{R}^d$ and $x \in \mathrm{conv}(X)$, there is an algorithm that runs in time $\beta(n,d)$ and gives a set $C \subset X$ of at most $d+1$ points such that $x \in \mathrm{conv}(C)$.*

Then, there is an algorithm that, given a set X of n points in $\mathbb{R}^d$, finds a Tverberg partition of size $\left\lceil \frac{\alpha n}{d} \right\rceil$ in time $O\left(f(\alpha, d, n) + \beta(n,d) \cdot \alpha n \cdot \mathrm{poly}(d)\right)$.

Proof of Lemma 2.2. Given a set X and a point z, denote by $\mathrm{depth}(z, X)$ the minimum number of points in any halfspace containing z. We begin by applying the first algorithm to find a point z such that $\mathrm{depth}(z, X) \geq \alpha n$, using time $f(\alpha, d, n)$. Then, as long as $\mathrm{depth}(z, X) > 0$, we can find a set C of at most $d+1$ points such that $z \in \mathrm{conv}(C)$, using the second algorithm. We may assume that if $|C| = d+1$, it is minimal, by checking if z lies in the convex hull of any d-element subset of C. Notice that for each closed halfspace H^+ such that its boundary hyperplane contains z, we have $H^+ \cap C \leq d$. Removing C from X reduces $\mathrm{depth}(z, X)$ by at most d. Thus, we can repeat this process at least $\frac{\alpha n}{d}$ times. The family of sets C obtained in each step gives us a Tverberg partition that intersects in z, as desired. □

Notice that if finding the point z employs an algorithm with a probability ε of failing, then the overall algorithm also has probability ε of failing.

Theorems 1.2 and 1.3 follow from the method above, given the algorithms for centerpoints presented in Clarkson et al. [9]. In order to obtain the results for the integer lattice, we generalize the methods of Clarkson et al. using a notion of centerpoints with integer coordinates. The existence of these can be settled by using Doignon's theorem, an integer counterpart to Helly's theorem which was later rediscovered twice with integer optimization in mind [6, 25].

Theorem 2.3 (Doignon's theorem [12])**.** *Let $\mathcal{F}$ be a finite family of convex sets in $\mathbb{R}^d$. If every 2^d sets in $\mathcal{F}$ have an integer point in common, then so does the entire family $\mathcal{F}$.*

Lemma 2.4. *Given any finite set S of integer points in $\mathbb{R}^d$, there is an integer point z such that every halfspace containing z has at least $\left\lceil \frac{|S|}{2^d} \right\rceil$ points of S.*

Proof of Lemma 2.4. Consider the family

$$\mathcal{F} = \left\{ \text{conv}(K) : K \subset S,\ |K| > \frac{(2^d-1)}{2^d}|S| \right\}.$$

By construction, $\mathcal{F}$ satisfies the conditions of Doignon's theorem. Let z be an integer point in $\cap\mathcal{F}$. If a halfspace containing z had fewer than $\frac{|S|}{2^d}$ points of S, it would mean that z would be separated by a hyperplane by a subset of S of cardinality strictly greater than $\frac{(2^d-1)}{2^d}|S|$, contradicting the fact that it lies in $\cap\mathcal{F}$. □

In order to construct integer centerpoints in a manner analogous to [9], we recall the notion of λ-samples.

Definition 2.5. *Given a set $S \subset \mathbb{R}^d$ with n points, we call $S' \subset S$ a λ-sample if for any halfspace H with $|H \cap S| \geq 4\lambda n$, we also have*

$$\frac{|H \cap S'|}{|S'|} \leq \frac{|H \cap S|}{n} + \lambda.$$

In particular, a β-center of a λ-sample of S is immediately a $(\beta-\lambda)$-center of S, for $\beta \geq 4\lambda$. In [9], it was shown that a random sample of sufficiently large size, but depending only on d, λ, ε, is a λ-sample with probability at least $1-\varepsilon$.

Thus, given a set of n integer points and $0 < \lambda \leq 2^{-(d+2)}$, we can find an integer $\left(\frac{1}{2^d} - \lambda\right)$-center with probability at least $1-\varepsilon$ in the following manner: First, we find a λ-sample S' of the set with probability $1-\varepsilon$. Then, finding an integer $\frac{1}{2^d}$-center for S' is an integer programming problem where the number of constraints is fixed (as d, λ, ε are fixed) so it can be solved in polynomial time. We know that a solution exists by Doignon's theorem, and the solution must be a $\left(\frac{1}{2^d} - \lambda\right)$-center of the original set, as desired.

3. Algorithms for geodetic Tverberg partitions

In this section we present algorithms for Tverberg partitions in the geodetic convexity spaces of certain graphs. The partition conjecture was already established for trees by Jamison [19]; our proof of Theorem 1.5 gives the optimal Tverberg number and shows how such a partition may be found algorithmically.

Proof of Theorem 1.5. Let p be an arbitrary vertex of the tree G. Suppose that removing p splits G into r connected components $G_1, G_2, \ldots, G_r$, and let $a_i = |G_i \cap U|$ for all i. We first show that we can find p such that $a_i \leq k$ for all i. If this is not the case, assume without loss of generality that $a_1 \geq k+1$. Notice that there is a vertex p' in G_1 which is adjacent to p, so we replace p by p'. By doing this, $G_2, G_3, \ldots, G_r$ and $\{p\}$ merge into one big component, but since they had altogether at most k points of U, they do not cause a problem. $G_1 \setminus \{p\}$ may be split into more than one component, but at most one of them has more than k points of S. If this is the case, call it G_1'. Since G_1' has fewer vertices than G_1, this process must end after at most $|G|$ iterations.

In order to find the partition, define $G_0 = \{p\}$. It suffices to notice that if we take two vertices of U in different components G_i, G_j, their convex hull must contain p. For any partition of a set of $2k$ vertices in which no part has more than k vertices, there is a way to pair up points such that no pair lies in the same part. Therefore, we obtain the desired result. The desired pairing can be also be found algorithmically. Simply take the two parts with the largest number of vertices G_1, G_2, pick two arbitrary vertices $v_1 \in G_1, v_2 \in G_2$, make them an edge of the matching and set them aside. Then, continue in the same way with the rest of the graph. This process ensures that at every stage no component has more than half of the remaining vertices, so we can continue until the graph is empty. □

In order to show that the value $2k$ is optimal, take any set of size $2k-1$ of leaves of a tree. Since any partition of them into k sets must include a singleton, there is no Tverberg partition into k parts. From this theorem, we may verify Eckhoff's partition conjecture in the case of geodetic convexity spaces on trees:

$$r_k = 2k \leq 3(k-1)+1 = (r_2-1)(k-1)+1.$$

One must beware of the special case in which the Radon number of the tree is less than four. However, this only happens with trees with exactly two leaves, reducing to the case of Radon on the real line. It should be stressed that the algorithm presented here is similar to those in the previous section. Indeed, our point p is the equivalent of a centerpoint for U, as the convex hull of any subset of U of cardinality greater than $|U|/2$ contains p. The same idea is extended to another family of graphs in Theorem 1.6. In this case instead of a single centerpoint, we look for a pair of vertices that in some sense serve jointly as centerpoints.

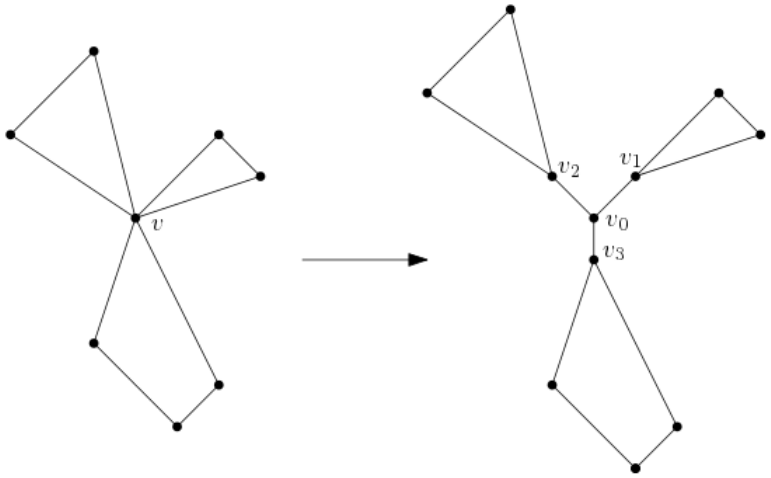

FIGURE 1. The construction of G' from G.

Proof of Theorem 1.6. If we could show that there are two vertices x, y such that no connected component of $G \setminus \{x, y\}$ has more than $2k-1$ vertices of U, we would be done. This is because in that case we could find a partition of U into $2k-1$ pairs such that the convex hull of each pair contains at least one of x, y. By the pigeonhole principle, this would imply that at least k pairs intersects in either x or y, yielding the partition we seek.

Notice that the condition of 2-separability means that two different cycles cannot share more than one vertex. Otherwise, we would be able to find two vertices with three edge-disjoint paths connecting them.

We construct an auxiliary graph G' as follows. First, we color the graph G blue. Then, consider any vertex v that is in more than one cycle. Notice that these cycles are pairwise edge-disjoint. If v is in n cycles, we replace v by n blue vertices $v_1, v_2, \ldots, v_n$ and set v_i adjacent to the vertices that v was adjacent to in the i-th cycle. Then, we include a red vertex v_0 which is connected to $v_1, v_2, \ldots, v_n$. An example is provided in Figure 1. We see that G' is connected and 2-separable. Moreover, in G' no two cycles share a vertex.

Let us show that no red vertex of G' is in a cycle. Suppose towards contradiction that v_0 is a red vertex which is in a cycle L_r. Notice that in the cycle v_0 must be adjacent to some v_i, which in turn is in a blue cycle L_b. If we follow L_r in the direction from v_i to v_0, at some point it must return to L_b. However, it cannot return to L_b through v_i, so the first time it gets back to L_b it must be using a different vertex x. Notice that this induces a path in G from v to x which does not use vertices of L_b, and contradicts the fact that G is 2-separable. Hence, no red vertex of G' can lie in a cycle.

Along with G', we define a copy U' of U. If v was a vertex of U in at most one cycle of G, we set its corresponding copy in G' to be part of U'. If v was a member of U in more than one cycle, we set only v_0 (the red

copy) to be part of U'. If we can find two vertices in G' such that removing them leaves only connected components with at most $2k-1$ points of U' each, the corresponding pair (or single vertex) in G would also work, as the only change could be that G would be shattered into more connected components.

Now we construct a second auxiliary graph G'' from G'. For each cycle in G', include a vertex in G'', and for every vertex in G' which is contained in no cycle, include a vertex in G''. Set u, v in G'' to be adjacent if there is an edge connecting some vertex corresponding to u and some vertex corresponding to v. Since G' is connected, so is G''. Let us show that G'' has no cycles.

If there were a cycle L' in G'', all the vertices of L' would correspond to cycles in G'. Take u one of its vertices. Using the cycle L', we can find a cycle in G' different from the one corresponding to u which intersects it. This contradicts the construction of G', in which no two cycles share a vertex.

Thus, G'' is a tree. Now to each vertex in G'' we assign the number of vertices in U' that correspond to it. Following the same method as in the proof of Theorem 1.5, we can find a vertex p in G'' the removal of which leaves connected components with at most $2k-1$ vertices of U' each. If p corresponds to a vertex in G', removing it (and any other point y) gives us the partition we seek.

If p corresponds to some cycle L in G', we have to do more work. For this, we can represent the vertices of L as the vertices of a regular $|L|$-gon in the plane in the order of the cycle. To each vertex v of L we assign two numbers. The first, $\mu_1(v)$ is 0 or 1 depending on whether v is in U' or not. The second, $\mu_2(v)$, is computed as follows. Notice that when we removed p from G'', we had several connected components left. Each of these corresponds to a connected component in G' which is connected to exactly one point of L. We take $\mu_2(v)$ to be the number of points of U' contained in connected components of $G' \setminus L$ which were connected to v in L.

Notice that μ_1 and μ_2 are two discrete finite measures in the plane whose sum is $4k-2$. By the discrete version of the ham sandwich theorem, there is a line ℓ which leaves at most half of each measure in each halfspace. We may assume without loss of generality that ℓ contains exactly two points x, y of L. Let us show that this choice of x, y satisfies the conditions we seek.

We let k_1, k_2 be the total measures of μ_1, μ_2, respectively, so that $k_1+k_2 = 4k-2$. When we remove the points x, y, we are left with at most two connected components which contain points of L. By the construction of ℓ, these have no more than $\lfloor k_1/2 \rfloor + \lfloor k_2/2 \rfloor$ points of U' each. We also have the connected components which were connected only to x or only to y, but these each have at most $2k-1$ points of U' by the construction of L. Thus, removing x, y gives us the partition we want, since each part contains at most $\max\{\lfloor k_1/2 \rfloor + \lfloor k_2/2 \rfloor, 2k-1\} = 2k-1$.

Notice as well that we may even allow for U to have repeated vertices, and the only change would be in allowing the vertices of G' to preserve such multiplicities in U', and in the definition of μ_2. □

4. Hardness of general Radon partitions

Proof of Theorem 1.7. Given a Boolean formula Φ in conjunctive normal form, we will describe a graph $G = (V, E)$ and $W \subset V$ with size poly(n), where a Radon partition (W^+, W^-) corresponds to a truth assignment satisfying Φ. This suffices to prove the theorem, as SAT is a known NP-hard problem, while counting the number of satisfying assignments is known to be #P-hard.

Claim. It suffices to consider Φ where every clause contains either all positive literals or all negative literals.

To prove the claim, suppose that Φ is given in terms of variables x_i. Replace every occurrence of $-x_i$ in Φ by a new variable y_i, and add the clauses $(x_i \vee y_i)$ and $(-x_i \vee -y_i)$ for every i, which ensures that x_i and y_i have opposite truth value. In the new formula Φ', every clause must have either all positive literals or all negative literals; and solutions to Φ' exactly correspond to solutions to Φ, as desired.

Suppose now that Φ takes the form given in the claim, with variables $x_1, \ldots, x_n$ occurring as positive literals in clauses $C_1^+, \ldots, C_\ell^+$, and occurring as negative literals in clauses $C_1^-, \ldots, C_m^-$. For convenience, we will also suppose that $\ell, m \geq 1$; if this is not true, we may simply add a new variable and a corresponding new clause containing it.

We now outline the construction of our graph G (shown in Figure 2). The vertex set V will consist of the following parts:

- the set $W^0 = \{w_1, \ldots, w_n\}$, corresponding to the variables $x_1, \ldots, x_n$,
- vertices w^+ and w^-, which will define which variables are True and False, respectively,
- the set V^+, which will allow us to calculate which of the clauses C_j^+ are True,
- the set V^-, which will allow us to calculate which clauses C_k^- are True.

The sets V^+, V^- overlap in the vertex v_0, which will allow us to test if all clauses C_j and D_k are satisfied.

The input W in the graph-Radon problem will be defined as $W^0 \cup \{w^+, w^-\}$. The motivation is that in any 2-partition of W such that w^+ and w^- are in different parts, we consider the variable x_i to be True if w_i is in the same part as w^+, and to be False if w_i is in the same part as w^-. We will eventually deal with the case where w^+ and w^- are in the same part; however, for the moment let us assume that each of the variables is designated unambiguously True or False.

Each vertex v in V, with the exception of v_0, possesses a parameter called *height*, denoted $h(v)$. The height of a vertex is a nonnegative integer, equal

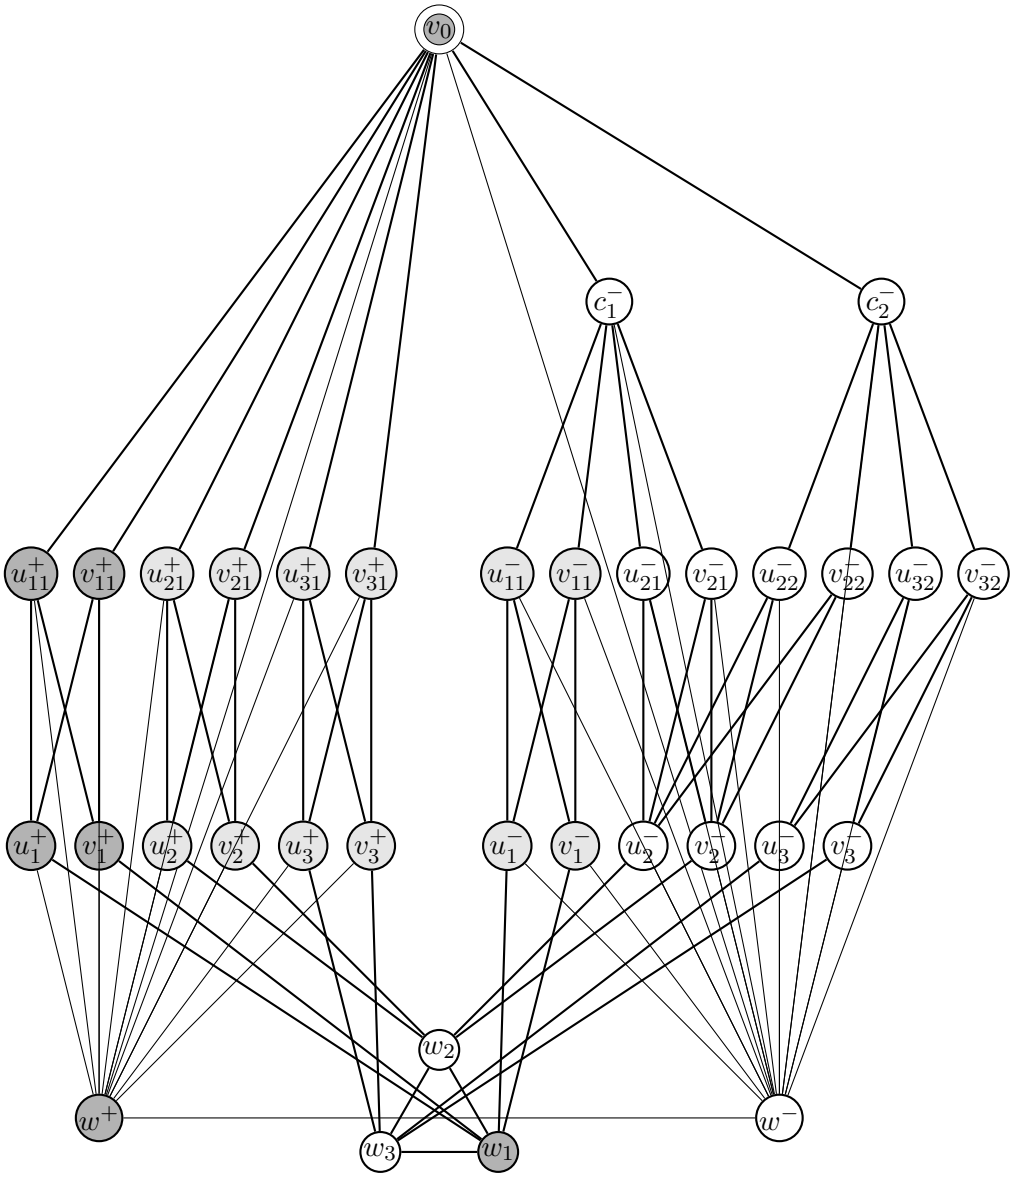

FIGURE 2. The graph G derived from the Boolean formula $\Phi = C_1^+ \wedge C_1^- \wedge C_2^-$, with $C_1^+ = x_1 \vee x_2 \vee x_3, C_1^- = -x_1 \vee -x_2, C_2^- = -x_2 \vee -x_3$. Note that the vertex v_0 could also be labeled c_1^+, a_1^+, and a_2^-. Darkly shaded vertices are in the convex hull of $\{w^+, w_1\}$, lightly shaded vertices are in the convex hull of $\{w^-, w_2, w_3\}$, and moderately shaded vertices are in neither. Note that vertex v_0 is in both convex hulls, indicating that the truth assignment $\{x_1 = \text{True}, x_2 = x_3 = \text{False}\}$ is a solution to Φ.

to 0 for $v \in W^0 \cup \{w^+, w^-\}$ and taking on various positive values for $v \in V^+$ or $v \in V^-$. The importance of the height is as follows: Suppose we start out with a subset of the height-0 vertices and wish to find the convex hull of these vertices. In each *extension step*, we add all vertices which lie on shortest paths between existing vertices. Then, we will design G such that in the first extension step we acquire only height-1 vertices, in the second extension step only height-2 vertices, and so on. We will never end up with more height-0 vertices than we started with, or with more height-t vertices than we acquired at extension step t. This means that we will be able to analyze the process of building the convex hull, one extension step at a time.

In addition to this property, we will construct G with the following property: The convex hull of a subset of W contains a vertex of $V^+ \setminus \{v_0\}$ only if it contains w^+; likewise it contains a vertex of $V^- \setminus \{v_0\}$ only if it contains w^-. Intuitively, V^+ is the part of the graph that handles True variables, and V^- is the part of the graph that handles False variables. We will now present the structure of V^+ and V^-. Given a partition of W in which w^+ and w^- lie in different parts, we will refer to the convex hull of the part with w^+ as the *positive convex hull*, and likewise define the *negative convex hull*.

At height 1 in G, we define vertices corresponding to two copies of each literal:

$$u_1^+, v_1^+, u_2^+, v_2^+, \ldots, u_n^+, v_n^+ \in V^+ \quad \text{and} \quad u_1^-, v_1^-, u_2^-, v_2^-, \ldots, u_n^-, v_n^- \in V^-.$$

The vertices u_i^+, v_i^+ are both included in the positive convex hull if (and only if) x_i is True, and u_i^-, v_i^- are both included in the negative convex hull if (and only if) x_i is False. To achieve this, we simply make u_i^+, v_i^+ adjacent to both w_i and w^+, while u_i^-, v_i^- are adjacent to both w_i and w^-. Thus, the shortest path between w_i and w^+ includes u_i^+, v_i^+, justifying our claim that height-1 vertices are included in the convex hull on the first extension step.

The construction of V^+ and V^- will be exactly symmetric from this point, and we will therefore present only the construction of V^+. The key here is that every vertex of V^+ is adjacent to w^+. The reason for this is that every shortest path between vertices of V^+ thus has length 1 or 2. It will be easy to ensure, therefore, that no pairs of vertices augment the convex hull in undesirable ways.

At height 2, we define vertices $u_{ij}^+, v_{ij}^+ \in V^+$ for each occurrence of literal x_i in clause C_j^+. (Likewise, we define $u_{ik}^-, v_{ik}^- \in V^-$ for each occurrence of literal $-x_i$ in clause C_k^-.) We let u_{ij}^+, v_{ij}^+ be adjacent to u_i^+ and v_i^+. Since u_{ij}^+, v_{ij}^+ lie in V^+, they are also adjacent to w^+. Then, u_{ij}^+, v_{ij}^+ lie on shortest paths between u_i^+ and v_i^+; thus u_{ij}^+, v_{ij}^+ are included in the positive convex hull at the second extension step if and only if x_i is True.

At height 3, we define a vertex $c_j^+ \in V^+$ for each clause C_j^+ such that c_j^+ is adjacent to u_{ij}^+, v_{ij}^+ for every value of i such that x_i occurs in C_j^+. Then, c_j^+ lies on a shortest path between u_{ij}^+ and v_{ij}^+; thus c_j^+ is included in the positive convex hull at the third extension step if and only if x_i is True. In

essence, we have defined an OR gate for the positive literals that appear in the clause C_j^+.

It may seem obscure why we defined u_{ij}^+, v_{ij}^+, instead of letting c_j^+ be adjacent to u_i^+, v_i^+ directly. The reason is that we must ensure that vertices of greater height do not cause vertices of lower height to be included in the convex hull. If c_j^+, d_j^+ were adjacent to u_i^+, v_i^+, then any literal x_i common to clauses C_j^+ and $C_{j'}^+$ would be switched to True whenever $c_j^+, c_{j'}^+$ were both included in the positive convex hull. This is undesirable; therefore, we introduce u_{ij}^+, v_{ij}^+, which provide additional space and prevent u_i^+ from being included in shortest paths from c_j^+ and $c_{j'}^+$. The reason for this is that any shortest path between c_j^+ and $c_{j'}^+$ must have length 2, since both vertices are adjacent to w^+.

Recall that we want to create a correspondence between solutions to Φ and Radon partitions of the vertices W. Given such a partition, corresponding to an assignment of True/False to every variable x_i, we have seen that each c_j^+ lies in the positive convex hull if and only if C_j^+ is satisfied. We will now define an AND gate to determine when all of the clauses c_j^+ lie in the positive convex hull.

For each j with $1 \leq j \leq \ell$, recursively define vertices $a_j^+ \in V^+$ at height $2 + j$ as follows: Let $a_1^+ = c_1^+$, and for $j \geq 1$ let a_j^+ be adjacent to a_{j-1}^+ and c_j^+. Observe that a_j^+ lies in the positive convex hull if and only if the vertices $c_1^+, c_2^+, \ldots, c_j^+$ lie in the positive convex hull. Likewise, we define vertices a_k^- for $1 \leq k \leq m$. The crucial detail is that we identify a_ℓ^+ and a_m^- as a common vertex v_0. (This is why the height of v_0 is not defined, since $2 + \ell$ may not be equal to $2 + m$.) Observe that for a partition of W, the vertex v_0 lies in both the positive and negative convex hulls exactly when every C_j^+ and every C_k^- is satisfied. Moreover, v_0 is the only possible overlap between the positive and negative convex hulls. Hence, Radon partitions of W correspond exactly to truth assignments that satisfy Φ.

One final detail remains: We assumed that, in any partition of W, the vertices w^+ and w^- lie in different parts. In order to enforce this requirement, we let W^0 form a clique, and add an edge between w^+ and w^-. For any partition of W in which w^+ and w^- lie in the same part, one of the two parts must contain only elements of W^0, and therefore forms a clique, which is its own convex hull. We have defined the graph G in such a way that the convex hull of $W^0 \cup \{w^+, w^-\} \backslash W'$ does not overlap W', for any subset $W' \subseteq W^0$. (The reason for the edge between w^+ and w^- is so that shortest paths between vertices of V^+ and V^- lie through w^+ and w^- and therefore do not intersect W'.) We conclude that, in any Radon partition, the vertices w^+ and w^- must indeed lie in different parts. This completes our proof. □

Our approach suggests an alternate proof for the statement (proven in [10]) that it is NP-hard to calculate the geodetic Radon number.

Proof of hardness of geodetic Radon number. Suppose that we are given a CNF formula Φ with variables $x_1, \dots, x_n$ and clauses $C_1, \dots, C_m$. Let us define a formula Φ^ℓ on variables $\cup_{j=1}^{\ell}\{x_1^j, \dots, x_n^j\}$. Let the clauses take the form $C_{i_1}^1 \vee C_{i_2}^2 \vee \cdots \vee C_{i_\ell}^\ell$ for all possible selections $i_1, \dots, i_\ell \in \{1, \dots, m\}$, where $C_{i_j}^j$ denotes the clause C_{i_j} written using the variables x_i^j. Thus, each clause of Φ^ℓ corresponds to an ℓ-tuple of clauses of Φ, and is satisfied if any of the constituent clauses is satisfied.

Claim 1. Φ^ℓ is satisfiable if and only if Φ is satisfiable.

If Φ has a satisfying assignment $x_i = t_i$, then we can create a satisfying assignment for Φ^ℓ by setting $x_i^j = t_i$ for each j. Conversely, suppose that Φ^ℓ has a satisfying assignment $x_i^j = t_i^j$. We claim that for some j, the assignment $x_i = t_i^j$ is satisfying for Φ. Suppose towards contradiction that for each j, there exists a clause $C_{i_j}^j$ of Φ which is not satisfied by the assignment $x_i = t_i^j$. Then, the clause $C_{i_1}^1 \vee C_{i_2}^2 \vee \cdots \vee C_{i_\ell}^\ell$ of Φ^ℓ is not satisfied by $x_i^j = t_i^j$, giving us our contradiction.

Claim 2. Suppose we remove up to $\ell - 1$ variables from Φ^ℓ (i.e., remove all occurrences of these variables without deleting clauses) to obtain $\tilde{\Phi}$. Then, $\tilde{\Phi}$ is satisfiable if Φ is satisfiable.

After removing up to $\ell - 1$ variables, there must remain some value j_0 such that none of $x_1^{j_0}, x_2^{j_0}, \dots, x_n^{j_0}$ is removed. If $x_i = t_i$ is a satisfying assignment for Φ, then assigning $x_i^{j_0} = t_i$ (and assigning all other x_i^j arbitrarily) gives us a satisfying assignment for Φ^ℓ, as desired.

Let $G(\Phi)$ denote the graph defined in the proof of Theorem 1.7, where Φ is a CNF formula. Now, consider the graph $G = G(\Phi^\ell)$, where we will pick ℓ to be a fixed constant, so that the size of G is polynomial in n. Let V^+, V^-, W^0, w^+, w^- be as in the preceding proof. Consider the problem of determining whether G has Radon number $\ell n + 2$. Let W be a subset of the vertices of G such that $|W| = \ell n + 2$. Note that if $W = W^0 \cup \{w^+, w^-\}$, then there exists a Radon partition of W if and only if Φ is satisfiable (by Claim 1 and Theorem 1.7). Hence, the Radon number of G is greater than $\ell n + 2$ if Φ is not satisfiable. We will show now that the Radon number is at most $\ell n + 2$ if Φ *is* satisfiable. We divide into two cases based upon the structure of W.

Case 1. Both $|W \cap V^+|$ and $|W \cap V^-|$ are at most $\lfloor (\ell - 1)/2 \rfloor$.

This condition implies that W differs from $W^0 \cup \{w^+, w^-\}$ in at most $\ell - 1$ vertices. Let us modify Φ^ℓ to $\tilde{\Phi}$ by removing the variables corresponding to these vertices. Then, Claim 2 implies that if Φ is satisfiable, then $\tilde{\Phi}$ is also; hence there exists a Radon partition of W.

Case 2. $|W \cap V^+|$ or $|W \cap V^-|$ is greater than $\lfloor (\ell - 1)/2 \rfloor$.

Without loss of generality, suppose $|W \cap V^+| \geq \lfloor (\ell - 1)/2 \rfloor$. Note that, by construction, there are no triangles within V^+; therefore, by choosing ℓ large enough, we can ensure that there exist vertices $y_1, y_2, z_1, z_2 \in W \cap V^+$ such that y_1, y_2 and z_1, z_2 are non-adjacent. Then, w^+ is included on shortest paths between y_1, y_2 and between z_1, z_2. Hence, any partition of W is a Radon partition if y_1, y_2 are in different parts, as are z_1, z_2.

We conclude that determining the Radon number of G allows us to infer the satisfiability of Φ, from which the result follows. □

Acknowledgments

The authors would like to thank Jesús De Loera and Reuben La Haye for their helpful comments during this work. D.R. was supported by a National Science Foundation Graduate Research Fellowship under Grant No. 1122374.

References

[1] I. Bárány and D. G. Larman. A colored version of Tverberg's theorem. *J. London Math. Soc.*, s2-45(2):314–320, April 1992.

[2] I. Bárány, S. B. Shlosman, and A. Szücs. On a topological generalization of a theorem of Tverberg. *J. London Math. Soc.*, 2(1):158–164, 1981.

[3] P. V. M. Blagojević, F. Frick, and G. M. Ziegler. Barycenters of polytope skeleta and counterexamples to the topological Tverberg conjecture, via constraints. *arXiv preprint arXiv:1510.07984*, 2015.

[4] P. V. M. Blagojević, B. Matschke, and G. M. Ziegler. Optimal bounds for a colorful Tverberg-Vrećica type problem. *Adv. Math.*, 226(6):5198–5215, 2011.

[5] P. V. M. Blagojević, B. Matschke, and G. M. Ziegler. Optimal bounds for the colored Tverberg problem. *J. Eur. Math. Soc. (JEMS)*, 17(4):739–754, October 2015.

[6] V. Borozan and G. Cornuéjols. Minimal valid inequalities for integer constraints. *Math. Oper. Res.*, 34(3):538–546, 2009.

[7] B. Bukh. Radon partitions in convexity spaces. *arXiv preprint arXiv:1009.2384*, 2010.

[8] C. Carathéodory. Über den Variabilitätsbereich der Koeffizienten von Potenzreihen, die gegebene Werte nicht annehmen. *Math. Ann.*, 64(1):95–115, 1907.

[9] K. L. Clarkson, D. Eppstein, G. L. Miller, C. Sturtivant, and S.-H. Teng. Approximating center points with iterative Radon points. *Internat. J. Comput. Geom. Appl.*, 6(03):357–377, 1996.

[10] E. M. M. Coelho, M. C. Dourado, and R. M. Sampaio. Inapproximability results for graph convexity parameters. *Theoret. Comput. Sci.*, 600:49–58, 2015.

[11] J. A. De Loera, R. N. La Haye, D. Rolnick, and P. Soberón. Quantitative Tverberg, Helly, & Carathéodory theorems. *arXiv preprint arXiv:1503.06116v3*, 2015.

[12] J. P. Doignon. Convexity in cristallographical lattices. *J. Geom.*, 3(1):71–85, March 1973.

[13] Mitre Costa Dourado, Dieter Rautenbach, Vinícius Gusmão Pereira de Sá, and Jayme Luiz Szwarcfiter. On the geodetic Radon number of grids. *Discrete Mathematics*, 313(1):111–121, 2013.

[14] Mitre Costa Dourado, Dieter Rautenbach, Vinícius Gusmão Pereira de Sá, and Jayme Luiz Szwarcfiter. Polynomial time algorithm for the Radon number of grids in the geodetic convexity. *Electronic Notes in Discrete Mathematics*, 44:371–376, 2013.

[15] Pierre Duchet. Discrete convexity: retractions, morphisms and the partition problem. In *Proceedings of the conference on graph connections, India*, pages 10–18, 1998.

[16] J. Eckhoff. Der Satz von Radon in Konvexen Produktstrukturen. I. *Monatsh. Math.*, 72:303–314, 1968.

[17] J. Eckhoff. Helly, Radon, and Carathéodory type theorems. In *Handbook of convex geometry, Vol. A, B*, pages 389–448. North-Holland, Amsterdam, 1993.
[18] J. Eckhoff. The partition conjecture. *Discrete Math.*, 221(1-3):61–78, 2000.
[19] R. Jamison. Partition numbers for trees and ordered sets. *Pacific J. Math.*, 96(1):115–140, September 1981.
[20] G. L. Miller and D. R. Sheehy. Approximate centerpoints with proofs. *Comput. Geom.*, 43(8):647–654, 2010.
[21] W. Mulzer and Y. Stein. Algorithms for tolerant tverberg partitions. *Internat. J. Comput. Geom. Appl.*, 24(04):261–273, 2014.
[22] W. Mulzer and D. Werner. Approximating tverberg points in linear time for any fixed dimension. *Discrete Comput. Geom.*, 50(2):520–535, 2013.
[23] S. Onn. On the geometry and computational complexity of Radon partitions in the integer lattice. *SIAM J. Discrete Math.*, 4(3):436–446, 1991.
[24] J. Radon. Mengen konvexer Körper, die einen gemeinsamen Punkt enthalten. *Math. Ann.*, 83(1):113–115, 1921.
[25] H. E. Scarf. An observation on the structure of production sets with indivisibilities. *Proc. Nat. Acad. Sci.*, 74(9):3637–3641, 1977.
[26] P. Soberón. Equal coefficients and tolerance in coloured Tverberg partitions. *Combinatorica*, 35(2):235–252, 2015.
[27] P. Soberón and R. Strausz. A generalisation of Tverberg's theorem. *Discrete Comput. Geom.*, 47:455–460, 2012.
[28] H. Tverberg. A generalization of Radon's theorem. *J. London Math. Soc.*, 41(1):123–128, 1966.
[29] M. L. J. van de Vel. *Theory of convex structures*, volume 50 of *North-Holland Mathematical Library*. North-Holland Publishing Co., Amsterdam, 1993.
[30] A.Y. Volovikov. On a topological generalization of the Tverberg theorem. *Math Notes*, 59(3):324–326, 1996.